

Consumption and Stock Prices: Can We Distinguish Signalling from Wealth Effects?

by

Nicolaas Groenewold,
Department of Economics,
University of Western Australia,
Crawley, WA 6009
Australia
Phone: +61 8 9380 3345
Fax: +61 8 9380 1016
e-mail: nic.groenewold@uwa.edu.au

* I am grateful to Patricia Fraser for updating figures for stock price fundamentals in Black et al. (2003) and for assistance with other data and to Gina Yoon for enthusiastic research assistance. Carol Bertaut has been generous in sharing her data, computing procedures and in answering many queries about the details of the work reported in her paper, Bertaut (2002).

Abstract

There has been a resurgence of interest in the effect of stock price changes on the real economy in the wake of the long stock market boom of the 1990s and the subsequent correction starting in 2000. One of the primary variables linking the stock market and output is consumption expenditure, with the wealth effect being the traditional channel of influence. More recently a number of other channels have been identified, in particular the signalling channel which sees stock prices as having simply a leading indicator effect. However, there has been little work which disentangles these channels empirically. This paper makes a contribution to this question by distinguishing between the effects of changes in stock prices driven by fundamentals and those driven by speculation. Since these two components of stock prices cannot be observed they must be generated by a model, and we use a decomposition recently applied by Black *et al.* (2003). Since any decomposition is likely to be controversial we experiment with various alternative decompositions. We find that both components of stock prices influence consumption but that the fundamental component is consistently the least important, thus supporting the wealth rather than the signalling channel.

JEL classification: E44

Key words: consumption, wealth effect, stock prices, stock price fundamentals

1. Introduction

The consumption function is the core of the Keynesian macro model in both its theoretical and econometric forms and it is therefore not surprising that it has been widely tested for various countries and time periods.¹ In recent years there has been a resurgence of interest in the role of stock prices or stock-market wealth in determining consumption behaviour. While early empirical work such as that by Bhatia (1972) and Peek (1983) included capital gains as part of income in the consumption function, recent empirical work has included stock-market wealth or stock prices as a wealth term. This recently renewed interest has followed and was no doubt stimulated by the long boom in stock prices in the 1990s as well as by a falling saving rate. Moreover, as stock prices faltered and then fell from the middle of 1999, there was understandable concern that consumption, aggregate demand and subsequently output would follow.

The traditional influence of stock prices on consumption is via a wealth effect; the standard theory of consumption makes consumption depend on (permanent) income and wealth with wealth being positively related, *inter alia*, to stock prices. However, as early as 1990, Romer (1990) suggested an alternative explanation for the decline in consumption following stock price falls during the Great Depression. She argued that, in contrast to the standard wealth effect, the stock market collapse resulted in increased consumer uncertainty which, in turn, caused them to defer expenditure on consumer durables.

More recently, Poterba and Samwick (1995), Poterba (2000) and Ludwig and Sløk (2002) have identified further channels. Poterba and Samwick distinguish between the wealth effect and a signalling effect where stock prices rise in

¹ See Hymans (1970) for an early evaluation of the consumption function in econometric models; for a recent and critical review of time-series consumption functions see Attanasio (1999).

expectation of output increases in the manner of a leading indicator, a purpose for which stock prices have often been used. Ludwig and Sløk set out two forms of the wealth effect – the realised and unrealised wealth effects (which are not distinguished in the standard consumption theory with no borrowing constraints) – as well as a liquidity-constraint effect and a stock-option value effect. As a fifth channel they list Romer’s consumer confidence effect.

Empirical work on the stock-price-consumption relationship has included cross-section, time-series and mixed (panel) studies. An example of a cross-section study is the one by Parker (1999) who directly addresses the question of whether the rise in stock prices in the 1990s could have contributed to the fall in the saving rate over the same period. Starr-McCluer (2002) used cross-section survey data in an attempt to disentangle Poterba and Samwick’s two channels and found some support for the traditional wealth effect, particularly for households with substantial stockholdings.

Time-series studies include the seminal paper by Poterba and Samwick (1995) and subsequent work by Ludvigson and Steindel (1999), Shirvani and Wilbratte (2000), who investigate the possibility of asymmetrical effects, Edison and Sløk (2002), who distinguish between “new” and “old economy” stocks, and Bertaut (2002) who examines the stock-price-consumption relationship for a number of industrial countries. An early study based on panel data is the one by Mankiw and Zeldes (1991) which focusses on the equity-premium puzzle which they examine separately for stockholding and non-stockholding consumers. More recently and more directly related to the interest of the present paper, Ludwig and Sløk (2002) and Case *et al.* (2001) use a combination of cross-section and time-series data, the latter to

investigate the relative strength of stock-market and housing wealth effects on consumption.

The general consensus is that there is an effect of changes in stock prices on consumption, that it is significant but relatively small (clearly weaker than the effects of change in the value of housing wealth and insufficient to explain all of the recent fall in the saving rate) and that the effect is probably asymmetrical. There has been little success in disentangling the various possible channels. Romer's (1990) work on consumption during the Great Depression depended mainly on the timing of consumption and stock price changes and on measures of stock-price uncertainty to suggest the relative importance of the uncertainty channel. Poterba and Samwick (1995) used disaggregated consumption and stock-holding data to argue that the wealth effect was likely to be small relative to the signalling effect. However, as various authors such as Ludwig and Sløk (2002) have pointed out, different channels have quite different implications for policy, making it important to sharpen our empirical knowledge of their relative importance.

This study contributes to progress in this direction by following a suggestion by Poterba and Samwick that distinguishing the wealth and the signalling effects requires an empirical identification of fundamental stock prices since only fundamentals are relevant for the signalling channel. Since fundamental and non-fundamental (we call them speculative) components of stock prices are not observable, they must be generated by a model and, given the state of modelling stock prices, the subsequent decomposition will not be uncontroversial, as Poterba (2000, p. 106) has pointed out. We use a method recently applied to US stock price data by Black *et al.* (2003) which is a macroeconomic application of the linearised dividend-discount model developed by Campbell and Shiller (1987, 1988, 19989). It identifies

fundamentals as depending on future output and discount rates and uses the model to impose restrictions on a vector autoregressive (VAR) model which is used to generate a series for fundamentals. We go on to explore the sensitivity of our results to this particular decomposition of stock prices by using three alternative, less sophisticated methods: a regression-based approach and two decompositions based on financial ratios, the price/earnings ratio and the dividend yield.

Whichever decomposition method is used, we use a two-stage approach in which the decomposition of stock prices is effected first and the two components are then used as regressors in a vector error-correction model (VECM) in which shocks to the two components are separately simulated.

The structure of the paper is as follows. The next section sets out the model and discusses the data used to estimate the model. Section 3 reports on the estimated model while the following section discusses the results of model simulation. Conclusions are presented in the last section.

2. The Model and the Data

The focus of the paper is on the analysis of the relationship between consumption and stock prices and we therefore begin with a simple model relating real consumption expenditure to real personal disposable income, an interest rate and stock prices. We also experimented in a limited way with other variables such as the unemployment rate (following Bertaut, 2002) but restrict reporting of our results to the core model. The form of the relationship used was in the spirit of the simple specification in recent papers by Ludvigson and Steindel (1999), Shirvani and Wilbratte (2000), Bertaut (2002) and Ludwig and Sløk (2002). As in most of the

recent literature, we analyse the relationship between these variables in a framework which accounts for the possibility of non-stationarity and cointegration.

After estimating the basic model, we go on to decompose the stock price variable into two components, one capturing fundamentals and the other non-fundamentals or the speculative component, as we call it. Since these two components can not be observed, they must be generated from a model. We use the output from a model by Black *et al.* (2003) but since such modelling is likely to be controversial, we experiment with several alternatives to assess the sensitivity of our results to the method of decomposition used. We therefore use four alternative approaches:

- (a) the decomposition by Black *et al.* (2003) based on the application to macroeconomic data of the Campbell and Shiller (1987, 1988, 1989) linearised version of the dividend-discount model. The data for the fundamental component is taken from their paper (updated to the end of 2002) and the speculative component is then simply the residual after subtracting the fundamental component from actual stock prices.
- (b) a regression based approach where stock prices are regressed on output and an interest rate variable chosen to represent the two main components of the dividend-discount model, output being reflected in profits and hence dividends and the interest rate influencing wealth-holders' discount rate. The explained part of stock prices is then taken as the fundamental component with the regression residuals capturing the speculative component.
- (c) an approach based on the assumption, implicit in much popular discussion of the stock market that in the long run the earnings/price ratio is constant. We compute the fundamentals as the level of stock prices which would have been

observed if the price/earnings ratio had been constant at its long-run average and the speculative component is then again the residual.

- (d) an approach based on a similar assumption that the dividend/price ratio is constant in the long run.

For these four alternatives, we need data for real consumption, real disposable income, real interest rates, the unemployment rate, real stock prices, real output, the Black et al. (2003) fundamental real stock price series, real earnings and real dividends. Data used and sources are as follows:

- real consumption: real personal consumption expenditures, 1996 dollars, seasonally adjusted; source: FRED data base, Federal Reserve Bank of St Louis.
- real disposable income: real personal disposable income, 1996 dollars, seasonally adjusted ; source: FRED data base, Federal Reserve Bank of St Louis.
- real interest rate: inflation-adjusted 3-month Treasury bill rate; source: Ibbotson Associates.
- unemployment rate: civilian unemployment rate, seasonally adjusted, average of monthly rates; source: FRED data base, Federal Reserve Bank of St Louis.
- real stock prices: inflation-adjusted S&P500; source: Ibbotson Associates.
- real output: real Gross Domestic Product; source: Ibbotson Associates.
- earnings: real earnings; source: “Irrational Exuberance” file, updated, from Robert Shiller’s web-site, <http://aida.econ.yale.edu/~shiller>.
- dividends: real dividends; source: “Irrational Exuberance” file, updated, from Robert Shiller’s web-site, <http://aida.econ.yale.edu/~shiller>.

We begin our analysis by checking for stationarity of the data to be used in the core equation. ADF statistics are reported in Table 1.

Table 1: Stationarity

Test	lc	ly	ls	u	r
No trend					
ADF(0)	-0.9386	-1.7342	-1.6343	-2.1148	-8.6220
ADF(1)	-0.9244	-1.7522	-1.6864	-4.2230	-6.2307
ADF(2)	-0.8217	-1.6867	-1.6477	-3.5513	-4.5975
ADF(3)	-0.8194	-1.6444	-1.6732	-3.1060	-3.9797
ADF(4)	-0.8693	-1.8846	-1.6978	-2.8274	-4.4959
Trend					
ADF(0)	-1.6584	-1.5222	-1.7707	-2.0642	-9.0850
ADF(1)	-1.6996	-1.4994	-1.9176	-4.4015	-6.5989
ADF(2)	-2.3411	-1.5646	-1.8546	-3.6872	-4.7327
ADF(3)	-2.3763	-1.5974	-1.8986	-3.1990	-4.0021
ADF(4)	-2.0096	-1.3005	-1.9408	-2.8912	-4.6592

Notes: lc , ly , and ls represent the logs of real consumption, real disposable income and real stock prices, u denotes the unemployment rate and r the real Treasury Bill rate. The 5% critical value for the ADF statistic in the case without trend is -2.8748 and for the case with trend it is -3.4314

The evidence shows that the three variables lc , ly and ls are clearly non-stationary irrespective of whether a trend is present in the “Dickey-Fuller equation” or the number of lags. Tests of stationarity for their first-differences (not reported) show that they are unambiguously $I(1)$. On the other hand, r is clearly $I(0)$, irrespective of the number of lags or the presence of a trend term. The results for the unemployment rate are less clear-cut – stationarity is sensitive to both the number of lags and the presence of a trend term. We, therefore, investigated the significance of the trend term and the number of lagged terms necessary to eliminate autocorrelation in the residuals of the testing equation. Both with and without trend, two lags are necessary

to eliminate autocorrelation from the residuals and in both these cases u is stationary. The trend term is not significant in any of the equations for u . Thus, we conclude that the logs of real consumption, real disposable income and real stock prices are all $I(1)$ and that the unemployment rate and the real Treasury Bill rate are both $I(0)$ and proceed to model accordingly.

We proceed to an analysis of the possibility of cointegration of lc , ly , and ls for which we use the Johansen procedure. A preliminary matter is the choice of the number of lags in the VAR used as a framework for the Johansen tests. Standard lag length selection criteria for a VAR in the first differences produce some conflicting indications but a model with three lags and no trend is free of autocorrelation in the individual equation residuals at 5%. At 3 lags both the Akaike and Schwartz criteria suggest shorter lags but tests for autocorrelation at lag 2 show some autocorrelation in the output equation. The trend proved not to be significant in the VAR in first differences so that we will include it in the VECM only via the cointegrating vector (if one exists). We therefore use a lag of three in the first differences which translates to a lag of four in the VECM if the variables are cointegrated.

Tests for cointegration in a model with 4 lags and trend only in the cointegrating vector produced clear evidence of one cointegrating vector whether the maximum-eigenvector or trace test was used as can be seen in the results reported in Table 2. The single cointegrating vector has plausible coefficients:

$$lc_t = 0.6594 ly_t + 0.0081 ls_t + 0.0029 trend_t$$

A test for the restriction that the trend coefficient is zero produces a prob value of 0.000.

Table 2: Cointegration

Maximum-eigenvalue test				
H ₀	H _A	Statistic	95% CV	90% CV
r = 0	r = 1	26.7235	25.42	23.1
r ≤ 1	r = 2	12.7248	19.22	17.18
r ≤ 2	r = 3	3.9253	12.39	10.55
Trace test				
r = 0	r ≥ 1	43.3736	42.34	39.34
r ≤ 1	r ≥ 2	16.6501	25.77	23.08
r ≤ 2	r = 3	3.9253	12.39	10.55

Note: r represents the number of cointegrating relationships between lc , ly and ls .

The cointegrating vector reported above is similar to that produced by a simple OLS regression:

$$lc_t = 1.9338 + 0.6959 ly_t + 0.0159 ls_t + 0.0026 trend_t$$

We also experimented with a model which included r as an endogenous I(0) variable. There was no change in conclusions regarding cointegration and the inclusion of r in the VECM showed that it was significant only in the equation for the first-difference in the log of stock prices but had little effect on the IRFs. We therefore proceed to use a model only with the three basic variables: lc , ly and ls .

3. Model Estimation

The estimated VECM is reported in Table 3. The explanatory power of the equations is modest, especially that for stock returns. It should be recalled, though, that the equations for consumption and disposable income are in first difference form so that low values for R^2 are not unexpected. Further, the low explanatory power of the stock-return equation is consistent with the Efficient Markets Hypothesis. The *ecm* term is significant in two of the three equations, confirming the cointegration conclusions derived earlier. The diagnostics for the equations show that they are all free from autocorrelation, and heteroskedasticity and that their functional form is appropriate.

In the consumption equation at least one of the lags of each of the explanatory variables is significant and in all cases the sum of the coefficients is of the right sign – positive for each variable. Of particular interest for the purpose of the present paper is the result that two of the three lagged stock return terms are significant and positive. The magnitude of the sum of the coefficients is similar to those reported for the US in Bertaut (2002) for a shorter sample period and a more elaborate model. Using only the consumption growth equation and ignoring the *ecm* term (taken into account below), we find that a 10 percentage-point sustained increase in stock returns increases consumption growth by approximately 0.57 percentage points. This compares to a figure of 0.64 reported by Bertaut (2002). We therefore have enough confidence in the results reported above to proceed with our main exercise of examining the consequences of decomposing stock returns into fundamental and speculative components.

Table 3: The VECM

Regressor	Equation		
	dLC	dLY	dLS
<i>const</i>	0.2371	-0.3239	1.7954
	[0.012]	[0.007]	[0.071]
$dLC1$	-0.0599	0.2454	0.5541
	[0.443]	[0.013]	[0.501]
$dLY1$	0.0804	-0.1193	-1.0901
	[0.173]	[0.110]	[0.080]
$dLS1$	0.0268	0.0254	0.0765
	[0.000]	[0.003]	[0.277]
$dLC2$	0.2659	0.0465	0.8331
	[0.000]	[0.625]	[0.294]
$dLY2$	0.1072	-0.0135	-0.7895
	[0.075]	[0.858]	[0.214]
$dLS2$	0.0173	0.0122	-0.0581
	[0.015]	[0.172]	[0.435]
$dLC3$	0.1150	0.0948	0.3030
	[0.108]	[0.294]	[0.687]
$dLY3$	-0.1202	-0.0672	-0.4243
	[0.034]	[0.345]	[0.475]
$dLS3$	-0.0051	-0.0063	0.0276
	[0.473]	[0.484]	[0.715]
<i>ecm1</i>	-0.0186	0.0265	-0.1430
	[0.014]	[0.006]	[0.074]
R^2	0.2682	0.1684	0.0436
AC	[0.704]	[0.183]	[0.340]
FF	[0.539]	[0.379]	[0.026]
Het	[0.360]	[0.016]	[0.919]

Notes: dLC_i is the i th lag of the first difference of the log of real consumption, dLY_i is the i th lag of the first difference of the log of real disposable income and dLS_i is the i th lag of the first difference of the log of real stock prices. The term “*ecm1*” is the first lag of the cointegrating regression residual. The numbers in brackets below coefficients are prob values for the hypothesis that the coefficient is zero. The numbers in the AC, FF and Het rows are also prob values – in the AC row for an LM test for first-to-fourth-order autocorrelation, in the FF row for the RESET test of functional form and in the Het row for an LM test of heteroskedasticity.

4. Model Simulation

We report simulation results in terms of impulse response function (IRFs) which capture the effects of a shock when all equations and the ecm terms are taken into account. In particular, we use the generalised IRFs devised by Koop *et al.* (1996) and Pesaran and Shin (1998). The nature of the IRF and the distinction between the generalised IRF and the more common form based on the Choleski decomposition is most conveniently exposted in a standard VAR model as follows.

Denoting the vector of variables in the VAR by $\mathbf{x}$, the model can be written as:

$$(1) \quad \mathbf{x}_t = \Phi(L)\mathbf{x}_t + \boldsymbol{\varepsilon}_t, \quad t = 1, 2, \dots, T$$

where $\Phi(L)$ is a p th-order matrix polynomial in the lag operator, L , where $L^n \mathbf{x}_t \equiv \mathbf{x}_{t-n}$.

We assume that $\mathbf{x}_t$ is stationary and that $E(\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t') = \Sigma$ is a positive definite matrix. We have ignored a constant (and other deterministic terms) for ease of exposition.

The IRF is easily derived from the vector moving-average (VMA) form of the model which can be obtained from (1) as:

$$(2) \quad \mathbf{x}_t = A(L)\boldsymbol{\varepsilon}_t,$$

where $A(L) = (I - \Phi(L))^{-1}$, an infinite-order matrix polynomial in L . One way of generating IRFs from (4) is to set one of the elements of $\boldsymbol{\varepsilon}_t$ at a non-zero value (usually its historical standard error) and all the others at zero and then trace the effects through successive values of $\mathbf{x}_t$. However, this ignores the fact that the elements of $\boldsymbol{\varepsilon}_t$ will generally be correlated so that historically a shock to one of the elements of $\boldsymbol{\varepsilon}_t$ will be associated with changes in other of its elements. A common method of overcoming this difficulty is to re-define the error terms to make them orthogonal so that they can be shocked independently. This is generally achieved by using the Choleski decomposition of the contemporaneous covariance matrix of the

errors, Σ . Since Σ is positive definite there exists a lower-triangular matrix (not necessarily unique), Q , such that

$$(3) \quad QQ' = \Sigma$$

The model can then be written in terms of the transformed errors, $\xi_t = Q^{-1} \varepsilon_t$, which are orthogonal. In this case the value of the IRF for the i th element of x following a shock to the j th error term n periods after the shock is given by

$$(4) \quad IRF_{ij}(n) = e_i' A_n Q e_j, \quad i, j = 1, 2, \dots, m; n = 0, 1, 2, \dots$$

where e_i is the i th unit vector and A_n is the n th matrix in the matrix polynomial $A(L)$.

While this is a popular procedure, it has the weakness that the orthogonalisation is not unique and the resulting IRFs are not unique but depend on the order in which the variables enter the model. The main alternative method, the generalised IRF, is to shock a particular error and then to shock all other errors in a way which preserves the historical relationship between them (or some other assumed correlations). Pesaran and Shin (1998) show that this involves computing the counterpart to (4) as:

$$(5) \quad IRF_{ij}^G(n) = \sigma_{jj}^{-1} e_i' A_n \Sigma e_j \quad i, j = 1, 2, \dots, m; n = 0, 1, 2, \dots$$

where σ_{jj} is the j th diagonal element of Σ . The advantage of the use of the generalised IRFs is that they are not affected by the ordering of the variables in the model. However, since the shocks in this case are not orthogonal, the IRFs cannot simply be added as they can in the conventional Choleski case. This is not usually a serious weakness since they are generally inspected one at a time.

Consider now the application of derivation of IRFs from the estimated VECM reported in section 3, beginning with the base case where no distinction is made between fundamental and speculative components of stock prices.

(a) *Base case*

We begin by reporting the impulse response function (IRF) for the effect of a stock price shock on consumption using the base model reported in section 3. With three equations, the VECM generates nine IRFs but, rather than report all nine, we focus on the effects on consumption of a one-standard-error shock to stock prices which is shown in Figure 1.




Figure 1: The IRF for log consumption following a stock price shock.

The effect of a one standard-error shock to the stock price equation error is clearly to increase consumption with the strength of the effect rising steadily over the first year before subsiding and reaching a steady state at a level of about 75% of peak after about four years. Taking into account the magnitude of the standard error of the dls equation (0.0796), the implied long-run elasticity of consumption with respect to real stock prices is approximately 0.05 which is similar to the magnitude of the long-run effect calculated from the dlc VECM equation in isolation, although in interpreting this result we should keep in mind that the generalised IRFs are based on simultaneous shocks to all equations.

Consider now the consequence of splitting up stock prices into fundamental and speculative components. As explained above, we focus on a decomposition based on recent work by Black *et al.* (2003) before proceeding to an assessment of the sensitivity of our results by using alternative decompositions.

(b) Consumption and the Black et al. (2003) decomposition of stock prices

Black *et al.* based their decomposition of stock prices on an application to macroeconomic data of a linearised version of the dividend-discount model due to Campbell and Shiller (see Campbell and Shiller, 1987, 1988, 1989). While Campbell and Shiller work with traditional financial inputs into the model such as dividends and earnings, Black *et al.* use real GDP and interest rates to derive an aggregate decomposition of the inflation adjusted S&P500 index for the period. They start with the dividend-discount model in the form:

$$S_t = E_t \sum_{i=0}^{\infty} \left(\frac{1}{\prod_{j=0}^i (1 + \rho_{t+j}^*)} \right) \cdot Q_{t+i} \quad (7)$$

where S denotes a real stock price index, $Q_t = \beta Y_t$, Y is real output, β is a scaling constant and ρ^* is the real discount rate. This equation is log-linearised, using the procedure of Campbell and Shiller, and written in terms of the log “price-dividend” ratio, $\pi_t = s_t - q_{t-1}$, where s and q are the logs of S and Q , as:

$$\pi_t = K + \sum_{j=0}^{\infty} \mu^j E_t \Delta q_{t+j} - \sum_{j=0}^{\infty} \mu^j E_t r_{t+j} \quad (8)$$

where K is a linearisation constant, r_{t+j} is the continuous discount rate and its expectation is interpreted as shareholders’ required return. In order to use (8) to generate a series for π_t^* , the price-output ratio implied by the model and from it the implied or fundamental stock price, s^* , Black *et al.* followed Campbell and Shiller

and used a VAR model to generate expected real output growth, Δq_t ; for the expected real discount rate they made three alternative assumptions: a constant discount rate, a discount rate with a time-varying risk premium and a discount rate incorporating a time-varying risk-free component. Since their results were not sensitive to the treatment of the discount rate, we use only the first and simplest of these. In this case the third term in (8) is incorporated into the constant K to obtain:

$$\pi_t = K' + \sum_{j=0}^{\infty} \mu^j E_t \Delta q_{t+j} \quad (8')$$

The VAR used to forecast output growth is a two-variable one. Define the vector $\mathbf{z}_t = (\pi_t, \Delta q_{t-1})'$; the VAR is then written as:

$$\mathbf{z}_{t+1} = \mathbf{A} \mathbf{z}_t + \boldsymbol{\varepsilon}_{t+1} \quad (9)$$

where $\mathbf{A}$ is a (2x2) matrix of coefficients and $\boldsymbol{\varepsilon}$ is a vector of error terms. Using the VAR for forecasting, implies that expected growth can be written as:

$$E_t \Delta q_{t+j} = \mathbf{e}_2' \mathbf{A}^j \mathbf{z}_t \quad (10)$$

where $\mathbf{e}_2 = (0, 1)'$. Hence the value of π_t generated by the combination of the present-value model and the forecasting assumptions (denoted π^*) is:

$$\pi_t^* = \frac{k-r}{1-\mu} + \mathbf{e}_2' \mathbf{A} (\mathbf{I} + \mu \mathbf{A} + \mu^2 \mathbf{A}^2 + \dots) \mathbf{z}_t = \frac{k-r}{1-\mu} + \mathbf{e}_2' \mathbf{A} (\mathbf{I} - \mu \mathbf{A})^{-1} \mathbf{z}_t \quad (11)$$

which is the equation they use to generate π^* and hence the fundamental stock price series once they have estimated the VAR coefficients and the constant K' . Finally, the (log of) fundamental stock prices is computed as:

$$s_t^* = \pi_t^* + q_{t-1} \quad (12)$$

Returning now to the empirical application to our consumption model, actual and fundamental log share prices derived from the Black *et al.* procedure, denoted ls and $lsfb$, are shown in Figure 2.




Figure 2: Actual and fundamental log stock prices: Black *et al.*(2003) approach


It shows long swings in stock prices relative to their fundamentals – for much of the 1950s and 1960s stock prices exceeded their fundamental level while they were below fundamentals for most of the 1970s and 1980s. The recent long boom in the 1990s had prices above fundamentals only for the second half of the decade.

The VECM equation for the first difference in log real consumption using this decomposition is:


$$\begin{aligned}
 dlc_t = & 0.2319 - 0.1019 dlc_{t-1} + 0.1258 dly_{t-1} + 0.0427 dlssb_{t-1} - 0.0929 dlsfb_{t-1} \\
 & (2.00) \quad (1.26) \quad (2.12) \quad (3.89) \quad (1.64) \\
 & + 0.3734 dlc_{t-2} + 0.1742 dly_{t-2} + 0.0503 dlssb_{t-2} - 0.1679 dlsfb_{t-2} \\
 & (4.43) \quad (2.73) \quad (4.13) \quad (2.95) \\
 & + 0.2768 dlc_{t-3} - 0.0408 dly_{t-3} + 0.0059 dlssb_{t-3} - 0.0721 dlsfb_{t-3} \\
 & (3.34) \quad (0.65) \quad (0.60) \quad (1.65) \\
 & - 0.0144 ecm_{t-1} \\
 & (1.96) \\
 R^2 = & 0.3148, AC\ prob = 0.029, FF\ prob = 0.387, Het\ prob = 0.665
 \end{aligned}$$

where absolute values of t-statistics are shown in parentheses. The explanatory power of the equation is noticeably improved as shown by the value of R^2 and the diagnostics show an absence of heteroskedasticity and functional form misspecification although there is evidence of autocorrelation at 5%. The coefficients of the two components of stock prices are of opposite sign; three are significant at 5% and a further two at 10%. We can test whether the restriction implicit in the use of total stock returns is consistent with the data by testing for the equality of the coefficients of the two components at each lag; not surprisingly, such a test results in the rejection of the restriction (a Wald statistic of 14.87 with a prob value of 0.002) and the same is true of a test of the restriction that coefficients at corresponding lags are equal but of opposite sign (a Wald statistic of 8.43 with a prob value of 0.038). Both sets of coefficients are jointly significant – prob values for the joint hypothesis that three coefficients are zero are 0.000 and 0.008 for *dlssb* and *dlsfb*. It appears, therefore, that the decomposition of stock prices into their two components improves the ability of the equation to explain consumption behaviour.

If we examine the consumption equation in isolation and ignore the ecm term, we find a positive long-run effect of a shock to the speculative component of stock returns and a negative effect of a shock to fundamentals, both of magnitudes much greater than in the base case. If we take the whole model into account using the generalised IRFs, we obtain the following time-profile for the effect on log real consumption of a one standard-error shock to the speculative and fundamental components, respectively:



(a) The effect of a speculative shock



(b) The effect of a fundamental shock

Figure 3: IRFs for lc using decomposed stock price: the Black *et al.* approach

The IRF for the speculative shock is very similar in shape but somewhat smaller in magnitude to that for the shock to total stock prices derived from the base case and pictured in Figure 1. This suggests that the effect in the base case is largely driven by the speculative component. The smaller magnitude of the effect is surprising, though, since the coefficients for the speculative component in the VECM equation for $d lc$ are consistently larger than for total returns in the base case. Moreover, the effect of a fundamental shock is in surprising contrast to that suggested by the signs of the estimated coefficients which are all negative. To interpret these

effects we need to keep in mind the nature of the generalised IRFs as well, of course, as the fact that the IRFs are for (log) levels whereas the equation is for the growth rate. The standard error for the speculative component is larger than that of its fundamental counterpart – almost fivefold in this case – and the equation residuals are positively correlated (with a correlation coefficient of 0.77). Recalling that the generalised IRFs are based on a shock to all errors that preserve historical correlations, we can see that the generalised IRF for a fundamental shock is dominated by the speculative effects which accounts, at least partly, for its unexpected sign. Indeed, if we break this relationship between the shocks by generating IRFs based on the Choleski orthogonalisation, we find that the effect of a speculative shock is similar to the generalised IRF while the effect of the fundamental shock is negative for approximately four years before turning slightly positive.

Thus we find that the nature of the IRF for a speculative shock does not depend on whether the generalised or orthogonalised version is used but that this choice does affect the shape of the IRF for a fundamental shock. Whichever we use, it is clear that the effect of the speculative component dominates that of the fundamental component and gives the distinctive shape to the IRF for consumption following a shock to total stock prices pictured in Figure 1.

We now proceed to an assessment of the sensitivity of our conclusions to the type of decomposition used by repeating the analysis with three alternative decompositions – starting with a regression-based approach.

(c) Consumption and a regression-based decomposition of stock prices

Consider now the effects of splitting up share prices into fundamental and speculative components using a regression approach common in the literature on

investment and stock prices.² In this approach stock prices are regressed on variables thought to capture the fundamentals according to a simple dividend-discount model in which real output drives profits and so dividends and an interest rate captures changes in investors' discount rate. This is similar to the model underlying the approach used by Black *et al.* discussed above but uses a simple regression rather than using a restricted VAR.

We therefore regress the log of real stock prices on four lags of the log of real GDP and the current value and four lags of the inflation-adjusted Treasury Bill rate.³ The equation explains about 59% of the variation in the log of real stock prices. The fitted values from this regression are taken as fundamentals and the residuals as the speculative component. The logs of actual and fitted stock prices, ls and $lsfr$, are pictured in Figure 4.




Figure 4: Actual and fundamental log real stock prices: regression approach

The general pattern of the fundamental component is similar to that derived from the more formal approach above although there appears to be more volatility in fundamentals in the present case.

² Blanchard *et al.* (1993) and Galeotti and Schiantarelli (1994) are examples of this approach.


³ We use the current value of the interest rate since interest rates are known instantaneously whereas real output is known only with a lag.

We include the fundamental and speculative components denoted by $lsfr$ and $lssr$ as separate variables in the VECM. We report only the consumption equation since that is our main focus. It is (with absolute values of t-statistics in parentheses):


$$\begin{aligned}
 dlc_t = & 0.2036 - 0.1055 dlc_{t-1} + 0.1194 dly_{t-1} + 0.0278 dlssr_{t-1} + 0.0373 dlsfr_{t-1} \\
 & (3.48) \quad (1.35) \quad (1.95) \quad (4.21) \quad (3.05) \\
 & +0.3371 dlc_{t-2} + 0.1413 dly_{t-2} + 0.0190 dlssr_{t-2} + 0.0115 dlsfr_{t-2} \\
 & (3.61) \quad (2.26) \quad (2.70) \quad (0.91) \\
 & +0.2416 dlc_{t-3} - 0.1107 dly_{t-3} - 0.0052 dlssr_{t-3} - 0.0035 dlsfr_{t-3} + 0.0253 ecm_{t-1} \\
 & (2.68) \quad (1.88) \quad (0.73) \quad (0.31) \quad (3.39) \\
 R^2 = & 0.3045, \quad AC \text{ prob} = 0.177, \quad FF \text{ prob} = 0.055, \quad Het \text{ prob} = 0.509.
 \end{aligned}$$

The value of R^2 has risen marginally compared to the base case, suggesting that the decomposition better describes the data although not as well as the Black *et al.* decomposition does. The diagnostics show an absence of autocorrelation, heteroskedasticity and misspecification of functional form and the *ecm* term is significant. In contrast to the results derived from the Black *et al.* decomposition, the coefficients of the two components of real stock returns are the same sign at each lag, signs which correspond to those of ls in the base case. There is also a remarkable similarity in magnitudes which suggests that there was little point in identifying the two components separately if this decomposition is used. A Wald test that the coefficients of corresponding lags are equal produced a prob value of 0.660 providing strong evidence against the decomposition. The sums of the coefficients are also similar – the long run effects of a sustained increase in the fundamental and speculative components of stock returns are 0.0790 and 0.0860 respectively compared to a figure of 0.057 for the base case. All this suggests that both components

significantly affect consumption and in very similar ways. This is not, however, borne out by the general IRFs which are pictured in Figure 5.



(a) Effect of a shock to the speculative component



(b) The effect of a shock to the fundamental component

Figure 5: IRFs for lc using decomposed stock prices: regression decomposition

Notwithstanding the similarity of the coefficients in the consumption equation, the effects of the two shocks are dramatically different once the entire model is taken into account. In order to interpret the effects, it is again important to recall the characteristics of the IRFs – they represent the effects of one standard-error shocks, they are *generalised* IRFs and they include the effects of the whole model, the ecm terms included. The first source of difference may be the differences in the variability

of the two stock-price components – the equation standard error for the fundamentals equation is 0.0411 while that for the speculative equation is more than twice as large at 0.0894. Second, the generalised IRFs impose historical correlations between the equation errors when formulating the shock. The correlation between the residuals of the two stock-price component equations is -0.4441 . Thus the negative effect of the fundamental shock, despite the overwhelmingly positive coefficients in the consumption equation, reflect the combined effect of the historical negative correlation between the two shocks and the fact that the speculative shock has a standard error more than twice that of the fundamental shock. If we use orthogonalised IRFs then the effect of the speculative shock is little changed (although a little larger) and the effect of the fundamental shock is also positive initially but considerably smaller than the effect of the speculative shock. Finally, there is the usual consideration when moving from an inspection of individual equation to an IRF which is that the latter takes into account all the interactions in the model – all three equations plus the cointegrating vector. But, whether we use the generalised or orthogonalised IRFs and whether we use the regression-based decomposition or the Black *et al* procedure, the broad implications are similar: the speculative or non-fundamental component of stock prices has a greater effect on consumption, pointing to the wealth-effect rather than the signalling explanation of the response of consumption to stock price changes.

We continue with our assessment of the robustness of this conclusion by considering two decompositions of stock prices based on commonly used financial ratios.

(d) Consumption and a price/earnings decomposition of stock price

The financial press, market analysts and many academics commonly discuss share valuation in terms of financial ratios and use measures such as the price-earnings (P/E) ratio, the dividend yield or dividend-price ratio (D/P) and the book-to-market (B/M) ratio to determine whether shares are underpriced or overpriced. The P/E ratio is perhaps the most commonly cited. The underlying assumption of such analysis is that there is some long-term value to which the ratio returns. This can be based on the Gordon model of share valuation and has also been the subject of empirical work. Thus, Campbell and Shiller (1998, 2001) show that periods of high P/E ratios tend to be followed by falls in share prices over the next 10 years. In this sense, the P/E ratio gives an indication of the fundamental value of shares and can be used to predict long run share market performance.⁴ The evidence also gives weight to arguments that low P/E ratios in the late 1970s indicated that the market was underpriced; and that the record high values in the late 1990s signalled an overvalued share market that was headed for a substantial correction.⁵

It has been argued recently that share-price behaviour of the 1990s has been unusual with the implication that there has perhaps been a shift in the long-term P/E ratio (amongst other things) due to structural changes in the economy. There has, however, been no convincing explanation of why this shift has occurred, with most analyses managing to explain only a part of the long bull run in the 1990s. Campbell and Shiller (2001, p.12) articulate this argument well by outlining a number of the ways the US economy has changed during the twentieth century. They note that throughout all past innovations in transport, communication and business, the P/E and

⁴ Blielberg (1989, p. 31) notes that this predictive power cannot be used to time the market because an overvalued market can 'do well' for a long time. Rather the P/E ratio can be used as a mild indicator of the likely changes in share returns over the next several years.

⁵ See for example Kopcke (1997), White (2000) and Cole, Helwege and Laster (1996) for further discussion of this argument.

D/P ratios have stayed remarkably stable. There is no reason to think that the ‘new economy’ warrants a permanent break from the long-term relationship between prices, earnings and dividends. Perhaps the most compelling evidence of this can be seen by observing the dramatic falls in the stock market since early 2000, bringing valuation ratios more in line with historical averages. Figure 6 shows the annual values for the P/E ratio for the S&P500 for our sample of 1947-2002 taken from Robert Shiller’s web-site.⁶




Figure 6. Quarterly P/E ratio, 1947(1)-2002(4)

This evidence shows that over a long period the P/E ratio fluctuated substantially over time, but tends to gravitate towards its mean of around 14 and even the historically high levels of the 1990s are subsiding.

Given these data and the unresolved nature of the debate, we use the assumed long-term constancy of the P/E ratio to estimate fundamental share prices. We have no information on the “normal” level of P/E but this is not a problem for the purposes of our empirical work. Denoting the normal level by k , we write the log of real fundamental stock prices using this definition as:

$$lsfpe_t = \log(k) + \log(earnings)_t$$

⁶ At <http://aida.econ.yale.edu/~shiller>

so that the value we choose for the normal level of P/E affects only the level of the log fundamentals and so will be immaterial for the VECM slope coefficients and the IRFs. We therefore set it at zero in our model estimation. Figure 7 plots the logs of real fundamental and actual share prices over the period based on this procedure.⁷ The graph supports many commonly held beliefs regarding share price behaviour over the last three decades. Firstly, we can see a period of undervaluation during the second half of the 1970s and the early 1980s. The 1990s, characterized by strong economic growth and an emerging ‘equity culture’, represents a decade of persistent overvaluation of the S&P 500 followed by the correction in the last two years of the sample.




Figure 7: Actual and fundamental real stock prices: P/E based decomposition


The VECM equation for $d\ln c$ based on this decomposition is:

⁷ The log of the mean of the P/E ratio has been added to $\ln sfpe$ (which is then labelled $\ln sfpem$) to maintain comparability with earlier diagrams.


$$\begin{aligned}
dlc_t = & 0.1915 - 0.0857 dlc_{t-1} + 0.0731 dly_{t-1} + 0.0272 dlsspe_{t-1} + 0.0567 dlsfpe_{t-1} \\
& (1.82) \quad (1.07) \quad (1.21) \quad (4.02) \quad (2.30) \\
& + 0.2612 dlc_{t-2} + 0.1132 dlt_{t-2} + 0.0180 dlsspe_{t-2} - 0.0366 dlsfpe_{t-2} \\
& (3.37) \quad (1.85) \quad (2.54) \quad (0.99) \\
& + 0.1294 dlc_{t-3} - 0.1138 dly_{t-3} - 0.0044 dlsspe_{t-3} + 0.0097 dlsfpe_{t-3} \\
& (1.75) \quad (1.99) \quad (0.62) \quad (0.40) \\
& - 0.0135 ecm_{t-1} \\
& (1.77)
\end{aligned}$$

$$R^2 = 0.2681, AC\ prob = 0.852, FF\ prob = 0.755, Het\ prob = 0.545.$$

The explanatory power of the equation for consumption growth is effectively the same as the base-case consumption equation which does not distinguish between the two components of stock prices, suggesting that the distinction is not useful. Two of the speculative terms are significant and positive with coefficients of the same magnitude as those in the base case. Only the first lag of the fundamental component is significant (and positive). A test of equality of the coefficients of corresponding lags produces a Wald statistic of 3.9109 with a prob value of 0.271. A test of the joint significance of the fundamental terms results in a prob value of 0.091 so that they are significant at 10% but not at 5%. The speculative terms, on the other hand, are highly significant with a prob value for a joint significance test of 0.000 (Wald = 24.4726). These results suggest that most of the explanatory power of real stock prices in the base-case equation come from the speculative component. This is borne out by the IRFs of shocks to the two equations. As was the case when we used the previous two decompositions, the effect of a speculative shock is substantial and of a similar shape to that of a stock price shock in the base case as shown in Figure 8(a) while the effect of a shock to the fundamental component is quite different and smaller.



(a) The effect of a speculative shock



(c) The effect of a fundamental shock

Figure 8: IRFs for lc using decomposed stock prices: the P/E approach

As in previous cases, the generalised IRF for the effect of a speculative shock is dominated by the correlated fundamental shock; if we use orthogonalised IRFs instead, the pattern and magnitude for the speculative shock are very similar while for the effect of the fundamental shock the pattern is preserved but the magnitude of the shock is reduced by about 50%. We can conclude, therefore, that the speculative component has an effect which is similar in shape and magnitude to that of total stock prices in the base case but that the effect of a speculative shock is small, especially when we use the orthogonalised IRF.

(e) *Consumption and a dividend/price decomposition of stock prices*

We now consider our final decomposition, that based on the argument that in the long run the dividend yield is constant. As in the case of a constant P/E ratio, we assume the existence of a long-run constant level although, again, the long-run level assumed affects only the intercept of the VECM and not the slope coefficients or the IRFs and so we set it at zero in our estimation of the model. The implied fundamental (log) stock price index is pictured with the actual log stock price index in Figure 9.




Figure 8: Actual and fundamental stock prices: the D/P decomposition


The overall shape is similar to that based on the P/E ratio although it is smoother. The

VECM equation based on this decomposition is:


$$\begin{aligned}
 dlc_t = & 0.2492 - 0.0404 dlc_{t-1} + 0.0543 dly_{t-1} + 0.0248 dlssdp_{t-1} + 0.0309 dlsfdp_{t-1} \\
 & (3.04) \quad (0.51) \quad (0.89) \quad (3.64) \quad (0.77) \\
 & + 0.2766 dlc_{t-2} + 0.0858 dly_{t-2} + 0.0162 dlssdp_{t-2} + 0.0158 dlsfdp_{t-2} \\
 & (3.63) \quad (1.39) \quad (2.27) \quad (0.33) \\
 & + 0.1255 dlc_{t-3} - 0.1297 dly_{t-3} - 0.0058 dlssdp_{t-3} + 0.0115 dlsfdp_{t-3} \\
 & (1.73) \quad (2.28) \quad (0.80) \quad (0.28) \\
 & - 0.0224 ecm_{t-1} \\
 & (2.97)
 \end{aligned}$$

$$R^2 = 0.2797, AC \text{ prob} = 0.758, FF \text{ prob} = 0.103, Het \text{ prob} = 0.474$$

As was the case for the previous equations, the diagnostics show that the equation is free from autocorrelation, heteroskedasticity and misspecified functional form. The value of R^2 is similar to that for the equation estimated using total stock returns. In this case fundamental stock returns are insignificant at all lags while the first two lags of the speculative component are significant and positive and of similar magnitude to the corresponding coefficients in the base-case equation suggesting, again, that the speculative component is the main driver of stock price effects on consumption. This is borne out by the IRFs pictured in Figure 9.



(a) The effect of a speculative shock



(b) The effect of a fundamental shock

Figure 8: IRFs for lc using decomposed stock prices: the D/P approach

The time-profile of the effect of a shock to the speculative component is again very similar in shape and magnitude to that for the effect of a shock to stock prices as a whole in the base model. The effect of a fundamentals shock, by contrast, is small and predominantly negative. The sign of the fundamental IRF is largely driven by the correlation with the speculative component captured by the generalised IRF – if we use orthogonalised IRFs instead, the picture for the speculative effect is little changed but the fundamentals shock has a predominantly positive effect although the shape of the IRF changes little. Again, we can conclude that shocks to the speculative component have a similar effect to those which follow a shock to total stock prices in the base case and that fundamentals have a relatively minor effect on consumption.

5. Conclusions

This paper has been concerned with the effect of stock price changes on consumption. Following the work of Poterba and Samwick (1995), the literature has identified two main channels through which the effect operates: the wealth effect and the signalling effect. In the first, causation is clearly from stock price rises to consumption when households simply spend (part of) their increased wealth resulting from increases in its market value. In the second, stock prices play the role of leading indicators and rise because underlying fundamentals rise and these stronger fundamentals are subsequently also reflected in rising consumption.

In this paper we have explored an approach which allows us to distinguish the importance of these two channels empirically: we have used a decomposition of stock prices into fundamental and non-fundamental components and argued that the signalling channel would imply that most of the explanatory power of stock prices in

a consumption equation would operate through fundamentals and that the wealth effect implies that the speculative component is at least as important from the point of view of consumption.

Since these two components are not observable, they must be model-generated. We focus on a procedure recently employed by Black *et al.* (2003) based on a linearised version of the dividend-discount model due to Campbell and Shiller (1987, 1988, 1989). However, since the decomposition method is not likely to be uncontroversial, we also explore the consequences of the use of a number of alternative methods for measuring the fundamental component in order to gauge the sensitivity of our conclusions to the method of decomposition. We use three alternative to the Black *et al.* approach, one based on regressing stock prices on fundamental variables and two based on financial ratios commonly used in the analysis of stock price movements.

Our findings can be summarised as follows.

- real consumption responds positively in both the short and long runs to a positive shock to total stock prices; the magnitude of the response is consistent with that found in earlier studies such as Bertaut (2002);
- real consumption responds to a positive shock to the speculative component of real stock price with a pattern and magnitude which is similar to its response to total stock prices, suggesting that most of the effect of total stock prices is driven by its speculative component; this conclusion is not sensitive to whether generalised or Choleski-orthogonalised IRFs are used; and
- the real consumption response to a shock to the fundamental component of stock prices is sensitive to the method of decomposition and to the type of IRF used but is in all cases smaller than the response to a speculative shock.

We conclude, therefore, that the wealth effect is alive and robust but that the signalling effect is fragile. These findings are consistent with those of Starr-McCluer (2002) who used cross-section survey data but conflict with the conclusions reached by Poterba and Samwick (1995) who concluded that wealth effects on consumption were likely to be small relative to signalling effects.

References

- Attanasio, O. P. (1999), "Consumption", Chapter 11 in J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics, Volume I*, Elsevier, 741-812.
- Bertaut, C. C. (2002), "Equity Prices, Household Wealth and Consumption Growth in Foreign Industrial Countries: Wealth Effects in the 1990s", International Finance Discussion Paper Number 724, Board of Governors of the Federal Reserve System.
- Bhatia, K. B. (1972), "Capital Gains and the Aggregate Consumption Function", *American Economic Review*, 62, 866-879.
- Black, A., P. Fraser and N. Groenewold (2003), "US Stock Prices and Macroeconomic Fundamentals", *International Review of Economics and Finance*, in press. A 2001 working paper version of the paper is available at <http://www.econs.ecel.uwa.edu.au/economics/Research/2001/DP%2001.08.pdf>
- Blanchard, O. J., C. Rhee and L. Summers (1993), 'The Stock Market, Profits and Investment' *Quarterly Journal of Economics*, 108, pp. 115-136.
- Blielberg, S. (1989), "How Little We Know", *Journal of Portfolio Management*, 16, 26-31.
- Campbell, J.Y. and R. J. Shiller. (1987). "Cointegration and Tests of Present Value Models", *Journal of Political Economy*, 95, 1062-1088.
- Campbell, J.Y. and R. J. Shiller (1988), "Stock Prices, Earnings and Expected Dividends" *Journal of Finance*, 43, 661-676.
- Campbell, J.Y. and R. J. Shiller (1989), "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factor", *Review of Financial Studies*, 1, 195-228.

- Campbell, J. Y. and R. J. Shiller (1998), "Valuation Ratios and the Long-Run Stock Market Outlook", *Journal of Portfolio Management*, 24, 11-26.
- Campbell, J. Y. and R. J. Shiller (2001), "Valuation Ratios and the Long-Run Stock Market Outlook: An Update", National Bureau of Economic Research Working Paper 8221.
- Case, K. E., J. M. Quigley and R. J. Shiller (2001), "Comparing Wealth Effects: The Stock Market versus the Housing Market", Wellesley College Working Paper 2001-09, Department of Economics, Wellesley College.
- Cole, K., J. Helwege and D. Laster (1996), "Stock Market Valuation Indicators: Is This Time Different?", *Financial Analysts Journal* May/June, 56-64.
- Edison, H. and T. Sløk (2002), "Stock Market Wealth Effects and the New Economy: A Cross-Country Study", *International Finance*, 5, 1-22.
- Galeotti, M. and F. Schiantarelli (1994), "Stock Market Volatility and Investment: Do only Fundamentals Matter?", *Economica* 61, 147-165.
- Hymans, S. H. (1970), "Consumption: New Data and Old Puzzles", *Brookings Papers on Economic Activity*, 1:1970, 117-126.
- Koop, G., M.H. Pesaran and S.M. Potter (1996), "Impulse Response Analysis in Nonlinear Multivariate Models", *Journal of Econometrics* 74, 119-147.
- Kopcke, R.W. (1997), "Are Stocks Overvalued?", *New England Economic Review* Sept/Oct, 21-40.
- Ludvigson, S. and C. Steindel (1999), "How Important is the Stock Market Effect on Consumption?", *FRBNY Economic Policy Review*, July 1999, 29-51.
- Ludwig, A. and T. Sløk (2002), "The Impact of Changes in Stock Prices and House Prices on Consumption in OECD Countries", IMF Working Paper WP/02/1, Research Department, International Monetary Fund.

- Mankiw, N. G. and S. P. Zeldes (1991), "The Consumption of Stockholders and Nonstockholders", *Journal of Financial Economics*, 29, 97-112.
- Parker, J. A. (1999), "Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate", *NBER Macroeconomics Annual*, 317-370.
- Peek, J. (1983), "Capital Gains and Personal Saving Behavior" *Journal of Money, Credit and Banking*, 15, 1-23.
- Pesaran, M.H., and Y. Shin (1998), "Generalised Impulse Response Analysis in Linear Multivariate Models", *Economics Letters* 58, 17-29.
- Poterba, J. M. (2000), "Stock Market Wealth and Consumption", *Journal of Economic Perspectives*, 14, 99-118.
- Poterba, J. S. and A. A Samwick (1995), "Stock Ownership Patterns, Stock Market Fluctuations and Consumption", *Brookings Papers on Economic Activity*, 2:1995, 295-357.
- Romer, C. D. (1990), "The Great Crash and the Onset of the Great Depression", *Quarterly Journal of Economics*, 105, 597-624.
- Shirvani, H. and B. Wilbratte (2000), "Does Consumption Respond More Strongly to Stock Market Declines than to Increases?", *International Economic Journal*, 14, 41-49.
- Starr-McCluer, M. (2002), "Stock Market Wealth and Consumer Spending", *Economic Inquiry*, 40, 69-79.
- White, C.B. (2000), "What P/E Ratio Will the US Stock Market Support?", *Financial Analyst's Journal* 56, 30-38.